

# Parameterized state variable equations for temporal properties

victor.yodaiken

December 2016

## 1 Introduction

Pnueli's temporal logic[1, 2] allows assertions about system state to be qualified by "when" properties must hold. The temporal logic qualifiers seem particularly well suited to specification of properties such as "every process waiting to execute will eventually execute" or "a currently executing processes must execute for at least  $t$  seconds or fail or request IO before it stops executing". In this note I am going to show how to get the same effect informally - not in the sense of "less rigorously", but in the sense of "in working mathematics, not formal logic". The treatment here is based on parametric equations described in [3] drawing on primitive recursive functions and automata products. The equational technique can be used to define operators that work like the temporal qualifiers, and approach extends to compositional and parallel systems because it uses a more structured model of system state. For example, if we have a property of well behaved processes that they will eventually engage in some input or output, and we have a property of operating systems that a process will always eventually execute, we should be able to show interaction between the two properties.

## 2 Background

When Pnueli and his colleagues developed Temporal Logic, it was widely believed that a mathematics for specifying programs should be constructed on a basis of some formal method or formal logic and so we might consider

a state to be a map assigning values to formal symbols. Given propositional symbols  $R_p$ , meaning the process with identifier  $p$  is ready to execute and  $X_p$  meaning the process with identifier  $p$  is currently executing, each state maps these assertions to true or false. If  $R_p$  is true and  $X_p$  is false, then process  $p$  is waiting to run. If the states represent states of a discrete state system, then they will be part of a state graph or transition graph. Pnueli's contribution was to use techniques from modal logic to quantify assertions over the transition graph. A temporal quantifier Always modifies a predicate so Always  $Q$  is satisfied by the current state if and only if  $Q$  is true for the current state and for every reachable state starting from the current state. Similarly Eventually  $Q$  is satisfied if and only if there is a bound  $n$  so that any path from the current state of length  $n$  or more must visit a state where  $Q$  is true. We can require:

$$\text{Always}(R_p \text{ implies Eventually } X_p) \quad (1)$$

so any state reachable from the current state which maps  $R_p$  to true must satisfy Eventually  $X_p$ . And if a state satisfies Eventually  $X_p$  we can suppose there is some  $n$  so that within  $n$  steps in the state graph we will encounter a state that maps  $X_p$  to true. The other basic temporal quantifiers are Next and Until, both with reasonably obvious interpretations (and all have some not so obvious complexities).

I'm going to describe an alternative but similarly motivated approach based on equations rather than formal logic and relying on a more compositional and structured notion of state. The treatment begins by considering what events can change state. Suppose we have a (possibly infinite) set of discrete events called the "alphabet" so that every discrete event that can cause the state to change is identified in the alphabet. A finite sequence of events drives the system from the initial state to some destination state. The null sequence leaves the system in its initial state. A sequence  $zq$  obtained by concatenating sequence  $q$  to sequence  $z$ , first drives the system to the state determined by  $z$  and then follows  $q$  to some later state. If it is clear from context that " $a$ " is an event, I will write  $za$  to indicate appending event " $a$ " to sequence  $z$  to advance the system state by a single step (if it's not clear from context, I will be explicit and write *Append*( $z, a$ )). I am assuming the effects of events on states are deterministic so each sequence determines a state. Many researchers have argued that we need to use non-deterministic methods for a number of reasons, but their concerns can be addressed through other means.

Now consider maps over finite event sequences. Suppose we have an equation  $y = f(\sigma)$  where  $\sigma$  is a free variable over finite event sequences. I will call variables like "y" that are defined by these types of equations: "state variables" and use the distinctive typeface as a reminder that they depend on the parameters  $\sigma$  and  $f$ . Think of a state as a collection of state variables that extract information from the event sequence. The same sequence of events can mean different things to different systems and so the state variables describing those systems would extract different information. Following what I take to be the essence of Pnueli's idea, the map "f" and the event sequence will usually be left as implicit parameters once we develop a few techniques shown below.

---

If "y" is a state variable, then  $y = f(\sigma)$  for some "f".

- $y(z)$  is the value of y in the state determined by sequence  $z$  since  $y = f(z)$ .
- $y(null)$  is the value of y in the initial state since  $y(null) = f(null)$ .
- $y(\sigma a)$  is the value of y in the state after  $a$  drives it from the current state since  $y(\sigma a) = f(\sigma a)$ .
- $y(\sigma z)$  is the value of y in the state after the sequence of events  $z$  drives the system from the current state since  $y(\sigma z) = f(\sigma z)$ .
- If  $z$  is a state variable  $y(z)$  is the value of y in the current state *inside the component* determined by  $z$ . Suppose  $z = g(\sigma)$  and  $y = f(\sigma)$ . Then  $y(z) = f(z) = f(g(\sigma))$ . This type of composition is sufficient for parallel composition as shown below.

---

To illustrate how this works in simple cases, suppose  $R_p$  and  $X_p$  are well defined state variables for each process  $p$ , so that  $R_p = g_p(\sigma)$  and  $X_p = h_p(\sigma)$  for some "g" and "h". And suppose we have some measure  $m : Events \rightarrow \mathcal{R}$  which tells us how much time passes during each event. Define  $L_p$  to sum up the elapsed time since the initial state. Set  $L_p(null) = 0$ . Set

$$L_p(\sigma a) = \begin{cases} L_p + m(a) & \text{if } R_p(\sigma a) \text{ and not } X_p(\sigma a) \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

If  $R_p$  and  $X_p$  are well defined, then  $L_p = f(\sigma)$  for some "f" that is well defined on every sequence. Now we could require that:

$$L_p < k \quad (3)$$

to make the liveness specification simpler and more concrete or

$$(\exists k > 0) L_p < k \quad (4)$$

to be less specific. Or consider

$$D_p(\sigma a) = \begin{cases} D_p(\sigma) + m(a) & \text{if } X_p \\ D_p(\sigma a) = 0 & \text{otherwise.} \end{cases} \quad (5)$$

Then:

$$X_p \text{ implies } (X_p(\sigma a) \text{ or } D_p > k \text{ or Not } R_p(\sigma a)) \quad (6)$$

tells us that a process that starts executing must run for at least k seconds or become unrunnable (e.g. by failing or by requesting a blocking operation).

### 3 Defining standard temporal quantifiers

Suppose  $Q = g(\sigma)$  is a boolean valued state variable. Note that because  $\sigma$  is always a free variable "always" is an implicit property of true assertions. That is, if  $Q$  is a boolean valued state variable and  $Q(\sigma)$  is a theorem,  $Q$  is true in every state. But let's make the temporal operators explicit. As a first approach:

$$\text{Next } Q \text{ iff (for all events } a) Q(\sigma a) \quad (7)$$

$$\text{Always } Q \text{ iff (for all finite event sequences } z) Q(\sigma z) \quad (8)$$

$$\begin{aligned} \text{Eventually } Q \text{ iff (there is some } n) \text{ so that} \\ \text{(for every } z \text{ of length } n) (z \text{ has a prefix } s) Q(\sigma s) \end{aligned} \quad (9)$$

The problem with these definitions is that  $Q(\sigma a)$  or  $Q(\sigma z)$  is not necessarily defined. That problem is a basic complexity for all the temporal quantifiers

(and sparked a large number of publications). What happens if the system gets to a state where some event occurs with undefined effect or where there is no further forward action possible? When we assert Eventually  $Q$  are we requiring that  $Q$  must become true or only that if the system can advance far enough  $Q$  will become true. For the former, we would do something like this:

$$\text{PossibleNext } Q \text{ iff } Q(\sigma a) \text{ is defined.} \quad (10)$$

$$\text{Next } Q \text{ iff } ((Q(\sigma a) \text{ whenever } Q(\sigma a) \text{ is defined}) \text{ and PossibleNext } Q). \quad (11)$$

$$\text{Always } Q \text{ iff } Q \text{ and Next Always } Q. \quad (12)$$

$$\text{Within } n \text{ } Q \text{ iff } Q \text{ or } (n > 0 \text{ and Next Within } n - 1 \text{ } Q). \quad (13)$$

$$\text{Eventually } Q \text{ iff exists } n, \text{ Within } n \text{ } Q. \quad (14)$$

$L_p < k$  is more precise than Eventually  $X_p$ . Suppose  $R_p$  becomes false before  $X_p$  becomes true - it's most likely that if the process becomes unrunnable we don't want to require that it eventually execute. Or suppose that there are events that do not correspond to the passage of time or that some events correspond to a longer duration than others. We can express these things using the temporal quantifiers, but less simply and, in my opinion, less clearly. Another problem with "eventually" is that the bound is too loose for practical systems. One could satisfy property 1 with an unrealizable system that keeps increasing the bound  $n$  - so that after  $n$  steps the process must execute, then after  $10n$ , then after  $100n$ . In practice, what would be more useful is a constant upper bound on the number of steps or even one bounded linearly by the number of processes or that is a function of system load.

## 4 Composition and concurrency

Suppose we have some state variables that represent the state of an individual process. We might have a boolean valued variable **INOUT** that is true if and only if the process is waiting on input or output. The property:

$$\text{Eventually INOUT} \quad (15)$$

requires that the process can't spin indefinitely on internal computations but must communicate at some point. Property 15 is a property of individual

processes and it would be good to relate it to properties 1 or 3 above but those are properties of the operating system and almost certainly involve completely different event alphabets. The solution is to define a relationship between the event sequence of the operating system and event sequences for each process  $p$ . We could have a sequence valued state variable  $u_p$  that depends on the sequence of events of the operating system and relates it to the induced sequence of events for the process component. Generally it is sensible to have the event sequence of each component be null when the enclosing system event sequence is null:

$$u_p(\text{null}) = \text{null}$$

It makes sense that a process must be executing to change state so:

$$\text{If Not } X_p \text{ then } u_p(\sigma a) = u_p(\sigma) \quad (16)$$

The sequence remains unchanged unless the process is currently executing.

$$\text{If } X_p \text{ then for some } b, u_p(\sigma a) = \text{Append}(u_p, b) \quad (17)$$

where we are appending "b" to the sequence  $u_p$  which is the same as  $u_p(\sigma)$  to indicate that when the enclosing system advances by event "a", the component  $p$  advances by "b". It follows that If  $u_p = z$  and  $X_p$  then  $u_p(\sigma a) = \text{Append}(z, b)$  for some b. Consider what  $\text{INOUT}(u_p)$  means. Since  $\text{INOUT}$  is a state variable, there is some  $f$  so that  $\text{INOUT} = f(\sigma)$ . And  $u_p$  is a state variable so there is some  $r_p$  with  $u_p = r_p(\sigma)$ . Thus  $\text{INOUT}(u_p) = f(u_p) = f(r_p(\sigma))$  and  $\text{INOUT}(u_p)$  is a function of  $u_p$  which is a function of  $\sigma$ . We are, thus, relating a per-process property,  $\text{INOUT}$ , to a property of the operating system by using  $u_p$  to define a relationship between events in the operating system and events in the component. Suppose

$$\text{Not } \text{INOUT}(u_p) \text{ implies } R_p \quad (18)$$

Then by 1

$$\text{Not } \text{INOUT}(u_p) \text{ implies Eventually } X_p \quad (19)$$

It follows by 17 that

$$\text{Not } \text{INOUT}(u_p) \text{ and } u_p = z \text{ implies Eventually}(u_p = zb \text{ for some } b) \quad (20)$$

If we do a little arithmetic on event sequence length we can now conclude

$$\text{Eventually } \text{INOUT}(u_p)$$

## 5 Conclusions and context

In an earlier paper ([3]), this work is connected to state machine theory and recursive functions. A map on finite strings is essentially equivalent to a "transducer", composite functions reflect products of state machines. The chaining form  $\mathbf{x} = \mathbf{y}(\mathbf{z})$  where  $\mathbf{z}$  is also a state variable so  $\mathbf{x} = f(\mathbf{z}) = f(g(\sigma))$  corresponds to a general product of automata where the inputs to components are functions of system input and the feedback - the outputs of other components. To define communicating components, we use simultaneous primitive recursion on the event sequences. Suppose events to processes can include events  $(\text{'send'}, p, v)$  for "transmit value  $v$  to process  $p$ " and  $(\text{'receive'}, p, v)$  for receive value  $v$  from process  $p$ . (I'm stealing the notation *'symbol* from Lisp to indicate that *'send* and *'receive* are just tokens or symbolic identifiers.) Let  $\mathbf{u}_p(\text{null}) = \text{null}$  and require that:

$\mathbf{u}_q(\sigma a) = \text{Append}(\mathbf{u}_q, (\text{'receive'}, v, p))$  if and only if  
 there is some  $p$  so that  $\text{OUT}(\mathbf{u}_p)$  and  $\text{PORT}(\mathbf{u}_p) = q$  and  $\text{VALUE}(\mathbf{u}_p) = v$   
 and  $X_q$  and  $X_p$   
 and  $\mathbf{u}_p(\sigma a) = \text{Append}(\mathbf{u}_p, (\text{'send'}, v, q))$

## References

- [1] Zohar Manna and Amir Pnueli. *Temporal Verification of Reactive Systems: Safety*. Springer-Verlag New York, Inc., New York, NY, USA, 1995.
- [2] Amir Pnueli. The temporal logic of programs. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science, SFCS '77*, pages 46–57, Washington, DC, USA, 1977. IEEE Computer Society.
- [3] Victor Yodaiken. State equations for discrete state systems. *CoRR*, abs/1608.01712, 2016.